

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 December 2001 (06.12.2001)

PCT

(10) International Publication Number
WO 01/93064 A1

(51) International Patent Classification⁷: G06F 15/163,
11/00

(21) International Application Number: PCT/IL01/00498

(22) International Filing Date: 30 May 2001 (30.05.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/207,336 30 May 2000 (30.05.2000) US
60/208,499 2 June 2000 (02.06.2000) US
60/208,558 2 June 2000 (02.06.2000) US
60/266,443 6 February 2001 (06.02.2001) US

(71) Applicant (for all designated States except US):
EMIKOLO NETWORKS INC. [US/US]; 300 East
42nd Street, New York, NY 10017 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): RAINES, Moshe

[IL/IL]; 2 Kehilat Venezia Street, Tel Aviv 69 400 (IL).
GOLDE, Ittai [IL/IL]; 5a Shmuel Sharira Street, Rishon
Lezion 75 381 (IL). REGEV, Oded [IL/IL]; 6 Nahal Sorek
Street, 75 246 Rishon Lezion (IL). REGEV, Amit [IL/IL];
6 Nahal Sorek Street, 75 246 Rishon Lezion (IL). EYAL,
Ori [IL/IL]; 24 Hasharon Street, 81 504 Yavne (IL).


(74) Agent: G. E. EHRlich (1995) LTD.; Bezael Street 28,
52521 Ramat Gan (IL).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK,
SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA,
ZW.

(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,

[Continued on next page]

(54) Title: DYNAMIC PEER-TO-PEER NETWORK CONTENT-SERVING



(57) Abstract: Apparatus for use in association with a network content provider (92) for providing dynamic mirroring of objects (94-112) of said content over said network in cooperation with network nodes activated for interaction with said apparatus by caching downloaded objects (94-112), the apparatus comprising: a listing unit (74) for preparing, dynamically updating and storing a listing of at least one object of said content and activated nodes (76) that have been sent said respective object, and a list extractor (72), associated with said listing unit for receiving download requests sent to said content provider and for extracting details from said listing of an object specified in said download request, placing said extracted details into a response for sending to a requestor, and for reporting to said listing unit on requests originating from activated nodes.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

Dynamic Peer-to-Peer Network Content-Serving

Related Applications

The present application is related to and claims priority from the following USA provisional applications, 60/266,443 dated 6th February 2001, 60/207,336 dated 30th May 2000, 60/208,588 dated 2nd June 2000, and 60/208,499 dated 2nd June 2000, the contents of each hereby being incorporated herein by reference.

Background--Field of invention

This invention relates to distribution of data content over a network and more particularly but not exclusively to distribution using peer to peer methodology of data content over the Internet.

Background--Description of Prior Art

Originally web content for a single site was served from a single computer, that was either connected to the Internet via a network shared with many other computers or by a relatively slow modem. This worked well enough as long as web content consisted mainly of text files and the number of "hits" a site received was relatively small. When sites user bases grew, and content consisted of more and more images (JPG, GIF), these sites had to buy larger pipes (ISDN, T1, T3) that were devoted only to serving the site. With multimedia, file sizes have

continued to grow, and site owners are forced to deploy multiple servers per site, and buy larger and larger pipes. Both of these options are typically slow and expensive. The use of larger pipes is often not enough to satisfy user demand for large, fast downloads. Very large sites have started deploying mirror servers in different parts of the world, on different sub-networks, to increase the speed at which users can download large files, by having their servers close (geographically, and in terms of the network) to their users. This is expensive both in terms of the hardware deployment, buying the necessary pipes, physical space for the servers to reside, and maintaining networks of computers that are far away from each other physically, but must be synchronized.

CDNs (Content Distribution Networks) attempt to relieve these problems by facilitating fast and reliable content delivery to users. Based on demand criteria, CDNs layer a network of technological intelligence across network providers, enabling quick and efficient data delivery. CDN service is based on the deployment of intelligent servers and other networking elements on the edge of various transport networks. These servers monitor networking traffic and served data to users. The networking elements of CDNs, usually servers loaded with routing control and caching software, create a unified platform across disparate networks for a site owner to deploy on. CDN usage thus reduces some of the cost to a site owner, spreading the cost of buying different machines for example, and also perhaps spreading the difficulties in maintaining them, over numerous site

owners. Nevertheless, as file size grows, CDNs have to buy larger and larger pipes, passing the cost onto their clients, those wishing to distribute content.

As user connections to various networks, including the internet, increase in terms of bandwidth, users continue to expect more and more content. CDNs do not help to solve the problem of the ever-increasing need for, and cost of, bandwidth to one wishing to deliver content to many different users.

As long as there have been computer networks, there has been the notion of file sharing. Operating systems allow files on a networked computer to be globally read. Recently there have emerged many different programs dedicated to allowing song, and other files to be searched for, and transferred across various networks including the internet. These systems still only allow for a single file at a time to be downloaded. The user has to manually search for a particular file. These systems cannot integrate many files from one computer together into a single experience for a user. These systems cannot integrate many files from separate computers into a single experience for a user. They force the user to pick which computer to transfer the file from. As it exists today, Peer-to-Peer (P2P) file sharing cannot help to solve the problem of those wishing to deliver content needing more and more bandwidth.

Summary of the invention

In accordance with one aspect of the present invention there is provided an application that caches web objects on different users machines, and redirects

other users to download those objects from users who have previously downloaded the same content object.

Accordingly, objects and advantages of the present embodiments include:

- (a) Reducing the bandwidth required by a content server.
- (b) Reducing the total bandwidth necessary to the entire network (Internet or Intranet)
- (c) Reducing the cost of administering a large web site
- (d) Reducing the physical cost of a large web site
- (e) Reducing the download time for users
- (f) Increases the reliability of a web site by
 - (1) Reducing the loads on the site's servers
 - (2) Providing multiple back ups in the event of total server failure.

According to one aspect of the present invention there is provided apparatus for use in association with a network content provider for providing dynamic mirroring of objects of said content over said network in cooperation with network nodes activated for interaction with said apparatus by caching downloaded objects, the apparatus comprising:

a listing unit for preparing, dynamically updating and storing a listing of at least one object of said content and activated nodes that have been sent said respective object, and

a list extractor, associated with said listing unit, for receiving download requests sent to said content provider and for extracting details from said listing of

an object specified in said download request, using said extracted details in a response for sending to a requestor, and for reporting to said listing unit on requests originating from activated nodes.

Preferably, said listing unit is operable to discriminate between activated sites on the basis of at least one of type of connection and bandwidth of connection, thereby to enable said response to guide requestors to easily accessible nodes.

An embodiment preferably further comprises an object expiration unit for applying to each object an expiration time for said activated nodes to remove said object from cached storage, said unit being associated with said listing unit, thereby to remove activated nodes whose objects have expired from said listing.

Preferably, the list extractor has network location functionality for determining relative locations between a requesting node and activated nodes having a requested object, thereby to favor relatively closeby activated nodes in said response.

Preferably, the apparatus has response tracking functionality for recording activated nodes to which extracted details have recently been sent, and operable to refer requesting sites to at least one of said activated nodes to obtain said details.

Preferably, the apparatus further comprises a node activation unit, said unit operable to send to a node a software agent able to cache objects and send said objects to other network nodes in response to requests therefor, thereby activating said node for use with said apparatus.

Preferably, the node activation unit is operable to determine whether a node is behind a firewall, said apparatus being directable not to regard nodes behind firewalls as activated nodes.

Preferably, the agent is operable to determine whether a node is behind a firewall and indicates said node as a non-activated node if it is behind a firewall.

Preferably, the agent comprises:

a request interception unit for intercepting requests for content from a user at said node to compare said request with a list of content sources associated with said apparatus,

request forwarding unit for forwarding said request to said apparatus,

request modifying unit for modifying said request to be directed at a node specified in said response.

Preferably, the agent is operable as an interface between said apparatus and a web-browser at said node.

Preferably, the request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

Preferably, the request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

Preferably the response comprises a referral to a listed activated node having a requested object so that the user is provided with the experience of

having received the object automatically and the entire referral process takes place without the user being aware of it.

According to another aspect of the present invention there is provided a network software agent, for use at a user node of a network in conjunction with dynamic mirroring apparatus, said agent comprising

a caching unit to cache internet data content objects, and

a cached object sending unit to send said objects to other network nodes in response to requests therefor,

thereby activating said node as a dynamic mirror for use with said apparatus.

Preferably, the agent is operable to determine whether a node is behind a firewall, said agent being directable not to activate said node behind a firewall as a dynamic node.

Preferably, the agent comprises:

a request interception unit for intercepting requests for content from a user at said node to compare said request with a list of content sources associated with said apparatus,

request forwarding unit for forwarding said request to said apparatus,

request modifying unit for modifying said request to be directed at a node specified in said response.

Preferably, the agent is operable as an interface between said apparatus and a web-browser at said node.

Preferably, said request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

Preferably, said request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

According to a further aspect of the present invention, there is provided a network software agent, for use at a user node of a network in conjunction with dynamic mirroring apparatus, the agent comprising:

a request interception unit for intercepting requests for content from a user at said user node to compare said request with a list of content sources associated with said apparatus,

request forwarding unit for forwarding said request to said apparatus,

request modifying unit for modifying said request to be directed at a dynamic mirror node specified in said response.

Preferably, the agent is operable as an interface between said apparatus and a web-browser at said node.

Preferably, said request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

Preferably, said request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

According to a further aspect of the present invention there is provided a method of dynamic mirroring of data content over a network comprising:

receiving a request for a content object,

checking for previous recipients of said content object,

sending to a requestor of said content a list of at least one previous recipient of said content, which said previous recipients are network nodes activated for caching of network objects for sending in response to redirected data requests, thereby dynamically mirroring said content objects over said network.

Preferably said previous recipient is supplied as an activated referral so that the entire process is invisible to the user. The automatic referral need not at the user end so that the user sends out the request via the user client, an ASM is provided including a favored alternative site and the object is downloaded directly to the user from the site recommended in the referral without the user client ever actually sending a request to the alternative site.

Brief description of the Drawings

In the drawing closely related figures have the same number but different alphabetic suffixes.

Fig. 1 shows a simplified schematic of the structure of a large network.

Fig 2 shows a server memory running a server application and a content serving site.

Figs 3a and 3b show a user memory before and after visiting a server operative according to an embodiment of the present invention.

Description of the Preferred Embodiments

A preferred embodiment of the invention requires a network, a content server having data content, typically arranged as objects, web pages, graphics files etc, and user machines on the network. A simplified view of such a network is given in Fig. 1. A content server 10 is connected via a pipe 12 to an ISP 14. The ISP has a pipe 16 running to a line owner 18, which is connected via pipes 20,22,24 to other ISP's 26,28,30 that in turn connect via various pipes 32-50 to end user computers 52-70. Computers 10,52-70 consist of memory, a CPU, and some mechanism to send and receive information from a network.

Reference is now made to Fig. 2, which is a simplified block diagram showing a server for providing data content over a network according to a first embodiment of the present invention. Computer 10 has content 94-112 to be distributed across a network (fig.1), content serving software 92, and server side software 72. On the server side application 72 there is a copy of a client side agent which end nodes would install on their systems 76, a common object list (COL) 74

which is a dynamically maintained list of what objects 94-112 have been downloaded by which network nodes or user computers 52-70.

The COL is a list of locations to which objects have been downloaded. The application may generate a subset of the COL, a personalized copy of the COL, an ASM which may contain references for the specific requesting end node to locations which a current end node may go to in the expectation of downloading the respective content object 94-112. The expectation is due to activating the locations to cache downloaded content for sending to subsequent users, as will be described in more detail below.

Reference is now made to Figs 3a and 3b, which are simplified block diagrams respectively showing an end user's machine before (3a) and after (3b) the user has visited a server application enabled content server. In fig. 3a, the user machine 52-70 has a content browser 114 connected via modem 116 or the like to the network 32-50. Upon visiting the server application enabled site, if the user hasn't already downloaded an agent program previously, either from the current site, or from some other location using the redirection technology, the user receives a small agent program 76. The agent program 76 sits between the browser and the network connection and acts as a proxy for the content browser. Proxy server is just one possible implementation, it may alternatively act as a packet sniffer, or connect directly to the users web browser, or even replace the users web browser, the agent may be a typical plug-in to the browser, or the agent may be a Java applet. The user agent 76 intercepts requests for content, checks

whether the request is to an content source that is part of the arrangement and if so, diverts the request to the server application 72, which replies by generating a subset of the COL, an ASM. The user agent downloads the necessary files for the browser 94,98,112, from one of the sites listed in the agents ASM and also caches them for sending on later to other users are redirected thereto. The currently requesting site is added to the COL with the information as to which objects it has downloaded so that future requests for the objects can be redirected thereto.

Preferably the address of the requesting site is used to determine at least an approximate physical location so that the site can be preferentially recommended to other sites closeby, and less to sites that are far away. Furthermore, in compiling the COL, as much information as possible is obtained as to the quality of the connection at the requesting site so that preferential status can be given to permanently connected sites with wide bandwidth connections as opposed to dialup nodes with narrow bandwidth connections.

Referring back to Fig.1, one skilled in the art may see how the present embodiments are not limited to networks of the size or specific structure of the Internet, but might also be used over a corporate Intranet, or any other network, with hubs, routers, and switches.

The appliance on the server may store each ASM, or may generate the ASM, or equivalents on the fly, based on the COL 74. The ASM's might be generated per machine, per sub-network, per ISP or any other network division.

The user agent may store an ASM on the client machine. It might connect each time to the server application to obtain a list of file locations. Depending on different features, any given user agent may not necessarily cache the content objects. For example, if the node is behind a firewall, it may be deemed pointless to cache the objects since requests for data downloads are likely to fail.

The dynamic mirroring application may reside on the content server, as in the above-described example or on a separate machine. In the latter case a single dynamic mirroring application may be applicable to numerous content servers. The separate application server might keep ASM's or generate them on the fly, might keep them per machine, per sub-network, per ISP, or any other network division, just as the content sever based application may.

The above-described embodiments are based on the assumption that content servers are sent request for the same content. Traditionally, content servers have responded to requests by downloading the same file, over and over. As user bandwidth has increased, and users expect greater content, content objects grow ever larger and it becomes more and more difficult for servers to handle the requests made by thousands if not millions of users. The system is designed to reduce the load on content servers, and reduce the overall network traffic.

User 52-70, connecting to an application enabled server 10 preferably has a thin (small) agent 76 installed on their machine, as described briefly above. The agent may act as a local proxy server or as a packet sniffer, or connect directly to the users web browser, or even replace the users web browser **fig.3** between the

user, and the content server. The agent caches some or all of the objects **fig.3** it downloads. The user agent keeps the server application updated as to which objects it has cached. A single agent will work with all servers that use the application.

The server application uses information from the user agent to build and maintain a dynamic Common Object List (COL) **74**. This is a list of users having each of the available objects cached on their machines. With the COL, the server application may create an individually tailored ASM, where the ASM is a machine specific (or sub net specific or ISP specific, or any other component of a network specific) list of other agents who have objects that a current agent is requesting. When a user with an agent tries to download an object, the agent intercepts the request. If that object belongs to an application enabled server, the agent checks its ASM **fig. 3, 78**, and if there is a copy of that object with another agent the agent downloads the requested object from the other user's agent. Alternately, if the agent has no ASM file for a given site, the user agent asks the server application for either an updated ASM for the site, or if there is another available agent with the requested object. The server application can then give the user agent a list of possible agents from the COL **74**. If any of the user agents obtained from the COL **74** already have the requested object, the requesting agent preferably downloads the requested object from an agent having the object, instead of having to download the object from the content server. If no other agent has the content object requested, the agent downloads the object from the content server. When a

second agent requests the same object, that second agent may be redirected to the agent already having the content object, instead of having to download it from the content server.

Once an object has been downloaded **fig. 3b, 94,98,112** it will rarely need to be downloaded from the content server again, as the next user can download the object from the first agent to download the object, and then there are two sources for a subsequent user to download from. As the user base grows, rarely will the content server have to upload a content object. The server application may maintain lists of where given objects are currently obtainable. Either the server application can upload entire lists to agents, or the server application may pass around references to objects, a task that is much lighter than having to serve out the same, large file over and over.

When there are many agents with many copies of the requested files, the server application may create ASM files for user agents. Those ASM files may contain those agents that are nearest (in terms of network distance) the agent the ASM file is being generated for. Being able to download files from locations close to the user may increase the download speed for the user. Downloading objects from near-by may also reduce the total network traffic by reducing the number of hops for any given content object.

The user agent is not limited to downloaded files from only one other agent at a time. The user agent may download one content object from one agent, and

another content object from a different agent. In this manner one agent can utilize its full pipe speed, even if the agents around it all have slower pipes.

The server application may mark each content object with a time-to-live marking. If the ASM is viewed as yet another content object, and the ASM's are generated per sub-network, instead of per machine, a machine with an ASM that is about to expire can look to other machines near it, and download a fresh ASM. In this manner, rarely will any object, including the ASM, need to be downloaded from the content server, or redirection application server, as the will only have to deploy each ASM file once per sub-net, ISP, or other network division.

It should be clear that by utilizing the bandwidth of end-users, a content redirecting server can dramatically reduce the load on content servers. By choosing agents that are close to one another, the server application/agent combination also reduces total network traffic. Again considering fig.1 and assuming that a user machine 52 requests an object, in order for that object to be downloaded from the content server 10, the object would ordinarily have to pass pipes 32, 20, 16, and 12. If a nearby user machine 54 has cached that object, instead of the object having to pass through four pipes, it may only have to pass through two pipes, 34, and 32. When the distances, and hops are greater and greater (a typically download may pass through more than 30 hops) it should be clear to one skilled in the art how reducing the number of hops content must pass through may enhance total network performance.

By reducing the bandwidth necessary for servers, the application may reduce, if not totally eliminate the need for conventional mirror servers. In this manner, the application may reduce hardware and maintenance costs.

When each user agent receives its own individual ASM, the agent may not have to talk to the content server, or server application at all. The user agent may well be able to assemble amalgamated content, including web sites from the content objects cached by other agents. When the agent does assemble amalgamated content including web sites from other agents, even in the case of a complete server failure, any end node which has an agent that can find a recently updated ASM or has an up to date ASM may still be able to view the content site.

The COL preferably keeps accurate lists of what has downloaded and by whom. The COL may thus be used not only to redirect agents to other agents, but also to compile accurate statistics about viewing. The COL keeps track of who got what from where, and who is near whom, therefore accurate records about network usage can be compiled. These records can help pre-generate ASM's to enhance efficiency, or the records might be used for any other purpose a site owner might wish.

Each user thus effectively becomes a mirror, and each agent may achieve faster download speeds even than can be managed with close mirrors. Instead of only being able to download an object from one source, the user will have many nearby sources to download the objects that compose the content from, and there is no single download bottle-neck. Even if the user has to go to the server

application for the ASM, the server application will only have to pass around object references, which is much easier than passing around full objects.

Furthermore the mirroring arrangement dynamically responds to user demand in that objects that are widely requested are automatically made widely available.

Accordingly the reader will see that using peer to peer technology in conjunction with records about network usage and which files have been downloaded by whom, along with an agent which can find and assemble all the necessary objects will greatly reduce the cost of maintaining a web site.

Traditionally, solutions such as deploying more and more bandwidth cost more and more as users expect greater content, and as users acquire greater bandwidth. By using peer-to-peer technology along with an agent that can seamlessly assemble content, the application/agent combination turns growing user expectation and bandwidth into a solution, instead of it being an ever growing problem. It is no longer necessary for a content server to get more bandwidth in order to keep up with ever increasing user demand. By harnessing the power of the growing user bandwidth, the application/agent combination enables faster user downloads with less network congestion. The application/agent combination enables greater server content without placing great additional new load on the server. The application/agent combination enables greater server reliability by creating many repositories of the server's content.

There are other configurations of the server application and agent. As has been noted, the agent may store the ASM file, and may cache content objects. For example, as mentioned above, if an agent is behind a firewall (and no other agent can see it, or download from it) then it need not cache any content objects. It may still have an active agent and retrieve objects from other agents, but would not need to store any. The server side application might pre-generate individually tailored ASM's and store them for users.

It should be clear that there would need to be security to prevent nodes (peers) from replacing cached objects with fake objects, prior to passing them on to other agents.. A preferred way to do prevent the substitution of fake data is to use a hash algorithm to calculate a hash code for each content object. These hash codes are then sent to the nodes as part of the nodes individually tailored COL, an ASM. User nodes compare the hash codes of any content that they receive to the expected hash codes (from the ASM) and can then see if the content was tampered with (on purpose or by accident).

There is preferably a management tool to optimize performance, altering the characteristics of the COL, setting it to store all objects, or just the largest objects, or only the most frequently requested objects, TTL (Time to live) of objects, TTL of ASM files, to recommend which agents to download from (someone who is downloading with many nodes nearby might only download from nodes on the same subnet, while someone who doesn't have as many agents near-by might download from anyone on his/her ISP) and any other variables.

It should be noted that the combination of server application and client agents may be used with dynamically generated content (CGI scripts, ASP, JSP, JHTML, XHTML, etc.). Thus a personally tailored version of a page would not be cached but a dynamically generated version of a page of interest to multiple users may be cached in its dynamically generated form.

It is appreciated that certain features of the invention, which are, for clarity, described in the context of separate embodiments, may also be provided in combination in a single embodiment. Conversely, various features of the invention which are, for brevity, described in the context of a single embodiment, may also be provided separately or in any suitable subcombination.

It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather the scope of the present invention is defined by the appended claims and includes both combinations and subcombinations of the various features described hereinabove as well as variations and modifications thereof which would occur to persons skilled in the art upon reading the foregoing description.

Claims

1. Apparatus for use in association with a network content provider for providing dynamic mirroring of objects of said content over said network in cooperation with network nodes activated for interaction with said apparatus by caching downloaded objects, the apparatus comprising:

a listing unit for preparing, dynamically updating and storing a listing of at least one object of said content and activated nodes that have been sent said respective object, and

a list extractor, associated with said listing unit, for receiving download requests sent to said content provider and for extracting details from said listing of an object specified in said download request, using said extracted details in a response for sending to a requestor, and for reporting to said listing unit on requests originating from activated nodes.

2. Apparatus according to claim 1, wherein said listing unit is operable to discriminate between activated sites on the basis of at least one of type of connection and bandwidth of connection, thereby to enable said response to guide requestors to easily accessible nodes.

3. Apparatus according to claim 1, further comprising an object expiration unit for applying to each object an expiration time for said activated nodes to remove said object from cached storage, said unit being associated with

said listing unit, thereby to remove activated nodes whose objects have expired from said listing.

4. Apparatus according to claim 1, said list extractor having network location functionality for determining relative locations between a requesting node and activated nodes having a requested object, thereby to favor relatively closeby activated nodes in said response.

5. Apparatus according to claim 1, having response tracking functionality for recording activated nodes to which extracted details have recently been sent, and operable to refer requesting sites to at least one of said activated nodes to obtain said details.

6. Apparatus according to claim 1, further comprising a node activation unit, said unit operable to send to a node a software agent able to cache objects and send said objects to other network nodes in response to requests therefor, thereby activating said node for use with said apparatus.

7. Apparatus according to claim 6, said node activation unit being operable to determine whether a node is behind a firewall, said apparatus being directable not to regard nodes behind firewalls as activated nodes.

8. Apparatus according to claim 6, said agent being operable to determine whether a node is behind a firewall and to indicate said node as a non-activated node if it is behind a firewall.

9. Apparatus according to claim 6, wherein said agent comprises:
a request interception unit for intercepting requests for content from a user at said node to compare said request with a list of content sources associated with said apparatus,
request forwarding unit for forwarding said request to said apparatus,
request modifying unit for modifying said request to be directed at a node specified in said response.

10. Apparatus according to claim 9, wherein said agent is operable as an interface between said apparatus and a web-browser at said node.

11. Apparatus according to claim 9, wherein said request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

12. Apparatus according to claim 9, wherein said request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

13. Apparatus according to claim 1, said response comprising a referral to a listed activated node having a requested object.

14. A network software agent, for use at a user node of a network in conjunction with dynamic mirroring apparatus, said agent comprising
a caching unit to cache internet data content objects, and
a cached object sending unit to send said objects to other network nodes in response to requests therefor,

thereby activating said node as a dynamic mirror for use with said apparatus.

15. A network software agent according to claim 14, being operable to determine whether a node is behind a firewall, said agent being directable not to activate said node behind a firewall as a dynamic node.

16. Agent according to claim 14, further comprising:
a request interception unit for intercepting requests for content from a user at said node to compare said request with a list of content sources associated with said apparatus,
request forwarding unit for forwarding said request to said apparatus,

request modifying unit for modifying said request to be directed at a node specified in said response.

17. Agent according to claim 16, operable as an interface between said apparatus and a web-browser at said node.

18. Agent according to claim 16, wherein said request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

19. Agent according to claim 16, wherein said request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

20. A network software agent, for use at a user node of a network in conjunction with dynamic mirroring apparatus, the agent comprising:

a request interception unit for intercepting requests for content from a user at said user node to compare said request with a list of content sources associated with said apparatus,

request forwarding unit for forwarding said request to said apparatus,

request modifying unit for modifying said request to be directed at a dynamic mirror node specified in said response.

21. Agent according to claim 20, operable as an interface between said apparatus and a web-browser at said node.

22. Agent according to claim 20, wherein said request modifying unit is operable to select an activated node from said response in accordance with proximity of said activated node over said network.

23. Agent according to claim 20, wherein said request modifying unit is operable to select an activated node from said response in accordance with a network connection quality of said activated node with said network.

24. A method of dynamic mirroring of data content over a network comprising:

receiving a request for a content object,

checking for previous recipients of said content object,

sending to a requestor of said content a list of at least one previous recipient of said content, which said previous recipients are network nodes activated for caching of network objects for sending in response to redirected data requests, thereby dynamically mirroring said content objects over said network.

25. A method according to claim 24, wherein said previous recipient is supplied as an activated referral.


26. A method of dynamic mirroring of data content over a network comprising:

receiving a request for a content object,

checking for previous recipients of said content object,

sending to a previous recipient of said content a request to forward to an originator of said request a copy of said object, wherein said previous recipients are network nodes activated for caching of network objects for sending in response to redirected data requests, thereby dynamically mirroring said content objects over said network.

1/3



SUBSTITUTE SHEET (RULE 26)




Fig. 2

3/3

Fig. 3a





Fig. 3b



INTERNATIONAL SEARCH REPORT

In tional application No.

PCT/IL01/00498

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/163, 11/00

US CL : 709/317; 714/114

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/317

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,611,069 A (MATOBA) 11 March 1997 (11.03.1997), col.9-col.10 and col.16-col.17	1-26
Y	US 5,790,789 A (SUAREZ) 4 August 1998 (04.08.1998), see Summary of the Invention.	1-26

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&"

document member of the same patent family

Date of the actual completion of the international search

21 September 2001 (21.09.2001)

Date of mailing of the international search report

06 NOV 2001

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

ALVIN E OBERLEY

Telephone No. 703 305-3665

Peggy Harrod

Form PCT/ISA/210 (second sheet) (July 1998)